

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
17 May 2001 (17.05.2001)

PCT

(10) International Publication Number  
**WO 01/35233 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 13/14**,  
13/00, 15/16, 15/00, 3/14, 17/30, H04L 12/28

John, A.; 3497 Cabrillo Ave., Santa Clara, CA 95051 (US). WEATHERFORD, William, T.; 1556 Forge Rd., San Mateo, CA 94402 (US).

(21) International Application Number: **PCT/US00/31243**

(74) Agent: HENNEMAN, Larry, E., Jr.; Henneman & Saunders, 121 E. 11th Street, Tracy, CA 95376 (US).

(22) International Filing Date:  
13 November 2000 (13.11.2000)

(81) Designated States (*national*): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(25) Filing Language: English

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(26) Publication Language: English

Published:  
— *With international search report.*


(30) Priority Data:  
09/440,016 12 November 1999 (12.11.1999) US

(71) Applicant: **AKAMBA CORPORATION [US/US]**; 15595 Los Gatos Blvd., Los Gatos, CA 95032 (US).

(72) Inventors: **SMITH, Jack, J.**; 28021 Audry Smith Lane, Saratoga, CA 95070 (US). **BURRIGHT, Richard, T.**; 2258 Dry Creek Ct., San Jose, CA 95124 (US). **WORLEY, W.**, Spencer III; 311 Correas Avenue, Half Moon Bay, CA 94019 (US). **MACDONELL, Eoin, B.**; 1057 Yorktown Dr., Sunnyvale, CA 94087 (US). **VASTANO,**

*[Continued on next page]*

(54) Title: **SYSTEM AND METHOD FOR MANAGING CONNECTIONS BETWEEN CLIENTS AND A SERVER WITH INDEPENDENT CONNECTION AND DATA BUFFERS**



**WO 01/35233 A1**



*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**SYSTEM AND METHOD FOR MANAGING CONNECTIONS BETWEEN CLIENTS  
AND A SERVER WITH INDEPENDENT CONNECTION AND DATA BUFFERS**

5

Jack J. Smith; Richard T. Burright; W. Spencer Worley, III;  
Eoin B. MacDonell; John A. Vastano; and William T. Weatherford

**RELATED APPLICATIONS**

This application is a continuation-in-part of co-pending U.S. Patent Application Serial  
10 No. 09/405,608, filed September 24, 1999, having at least one inventor in common herewith,  
and being under obligation of assignment to a common assignee. The parent application is  
incorporated herein by reference in its entirety.

**BACKGROUND OF THE INVENTION**

15 **Field of the Invention**

This invention relates generally to network servers and more particularly to servers  
that host a large number of client connections. Even more particularly, the present invention  
relates to servers (e.g., internet web servers) which host a large number of relatively slow  
client connections.

20 **Description of the Background Art**

It is common for network file servers such as internet web servers to host a  
large number of relatively slow client connections. The large number of open connections  
places a substantial burden on the server central processing unit (CPU), just to manage the  
open connections. For example, managing the open connections on a loaded server can  
25 consume 30-40 % or more of the CPU's operating capacity. This burden substantially  
reduces the percentage of CPU cycles available to perform the primary function of the server,  
i.e., providing data to clients.

The connection management burden on the server CPU degrades the performance of  
the server software routines and reduces the maximum number of client connections that can  
30 be open at one time. As a result, web-hosting companies must provide additional, redundant

servers to serve an increased number of clients. The cost of acquiring and maintaining additional web servers is substantial.

Proxy servers perform some client connection management functions, and are known in the art. However, it is well-known and commonly accepted in the art that such proxy servers must be housed separately from the server, and thus must communicate with the server over relatively slow, error prone network connections which the server must manage. See for example, Ari Luotonen, *Web Proxy Servers* (Prentice Hall, 1997), which is incorporated herein by reference.

What is needed, therefore, is a system and method for relieving the server CPU of the connection management burden, thus allowing the server to more efficiently host an increased number of clients.

### SUMMARY

The present invention overcomes the problems associated with the prior art by providing a system and method for managing connections between a plurality of clients and a server. The invention facilitates off-loading the connection management burden from the host CPU to an adapter card interposed between the network and the host bus.

The adapter card includes a network controller, a memory device, a processing unit, and a protocol adapter. The memory device provides storage for data and code. The code includes a proxy application that communicates with clients on the network via the network controller, and communicates with the server via the protocol adapter, which is coupled directly to the server bus.

When executed by the processing unit, the proxy application manages client connections by establishing network connections between the proxy application and clients via the network, and by establishing bus connections between the proxy application and the server via the server bus. Additionally, the memory device provides data buffering, which allows many network connections to be open with clients, while a relatively few bus connections are open to the server. In a particular embodiment, the proxy accumulates client data in the buffers from the large number of slow client connections, and then submits the client data to the server over the fast bus connections. Conversely, the proxy receives server

data via the fast bus connections, temporarily stores the server data, and then forwards the server data to the clients via the slow client connections.

In a more particular embodiment, the code includes a communications protocol stack that is employed by the application proxy to communicate with the clients and the server. In 5 an even more particular embodiment, the communications protocol stack is a Transmission Control Protocol/Internet Protocol (TCP/IP) stack.

In one embodiment, the server connections are opened only after the proxy determines that a complete client request has been received. The server connections are then closed after the proxy receives a response to the client request from the server. Optionally, a 10 predetermined number of persistent server connections are opened at system start-up, and the proxy uses these persistent connections to communicate with the server.

The proxy application optionally includes a number of application specific proxies, including but not limited to an HTTP proxy, a security proxy, and/or a pass-through proxy. In a particular embodiment, a master process module of the proxy discerns an application 15 identifier (e.g., a well known port number) from the client data, and invokes one or more of the application specific proxies corresponding to the value of the identifier.

A system and method for allocating buffers is also disclosed, whereby the number of client connections that can be opened and managed by a proxy application is substantially increased. According to the method, buffers are allocated to a client connection only after it 20 is determined that the buffer is needed to facilitate data transfer between a client and a server. One particular method includes the steps of establishing a connection with a client on behalf of a server, receiving a communication from the client, determining from the communication whether data will be exchanged between the client and the server (e.g., does the communication include a data request?), and allocating an input buffer to the client connection only if data is to be exchanged between the client and the server. Another 25 particular method includes the steps of establishing a connection with a client on behalf of a server, receiving a communication from the client, determining whether data will be received from the server (e.g., whether a complete data request has been received from the client), and allocating an output buffer only if data is expected from the server. In an exemplary embodiment, the methods of the present invention are implemented in an adapter card for 30 coupling a server to a network.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described with reference to the following drawings, wherein like reference numbers denote substantially similar elements:

5 FIG. 1 is a block diagram of a server and an adapter card according to the present invention;

FIG. 2 is a block diagram of the working memory of the adapter card of FIG. 1, showing the proxy module in greater detail;

10 FIG. 3 is a block diagram showing the application proxies module of FIG. 2 in greater detail;

FIG. 4 is a block diagram showing exemplary data structures for at least some of the data stored in the data buffer of FIG. 2;

FIG. 5 is a flow chart summarizing one method for managing connections between a client and a server according to the present invention;

15 FIG. 6 is a flow chart summarizing one method for performing the first step of the method of FIG. 5;

FIG. 7 is a flow chart summarizing one method for performing the second step of the method of FIG. 5;

20 FIG. 8 is a flow chart summarizing one method for performing the third step of the method of FIG. 5;

FIG. 9 is a flow chart summarizing one method for performing the fourth step of the method of FIG. 5;

FIG. 10 is a flow chart summarizing one method for performing the fifth step of the method of FIG. 5; and

25 FIG. 11 is a flow chart summarizing one method for performing the sixth step of the method of FIG. 5.

FIG. 12 is a block diagram showing an alternate client data structure and buffering scheme for at least some of the data stored in the data buffer of FIG. 2;

30 FIG. 13 is a block diagram showing the buffer status information of FIG. 12 in greater detail;

FIG. 14 is a flow chart summarizing an alternate method for performing the second step of the method of FIG. 5;

FIG. 15 is a flow chart summarizing an alternate method for performing the fourth step of the method of FIG. 5;

5 FIG. 16 is a flowchart summarizing a method for transferring data into an input buffer or an output buffer according to the present invention; and

FIG. 17 is a flowchart summarizing a method for transferring data out of an input buffer or an output buffer according to the present invention.

10

#### DETAILED DESCRIPTION

The present invention overcomes the problems associated with the prior art, by off-loading much of the connection management burden from the server's main processor with a proxy application run on a different processing unit. In the following description, numerous specific details are set forth (e.g., particular communications protocols, particular software and data structures, etc.) in order to provide a thorough understanding of the invention.

15 Those skilled in the art will recognize, however, that the invention may be practiced apart from these specific details. In other instances, details of well known network components and programming practices (e.g., establishing connections via a communications protocol stack) have been omitted, so as not to unnecessarily obscure the present invention.

20 FIG. 1 is a block diagram showing a system 100 coupled to an internetwork 102 via a physical network media 104. In a particular implementation, system 100 is an internet web server, and internetwork 102 is the Internet, but those skilled in the art will recognize that the present invention may be implemented in any type of network server.

System 100 includes a file server (e.g., an HTTP web server) 106 and an adapter card 25 108. File server 106 provides data to and receives data from clients 109(1-n) on internetwork 102, via adapter card 108. Adapter card 108 establishes and maintains network connections between clients 109(1-n) and adapter card 108, and establishes bus connections between server 106 and adapter card 108. Thus connected, adapter card 108 receives communications from clients 109(1-n) on behalf of server 106, forwards the communications to server 106, 30 receives responses from server 106 on behalf of clients 109, and forwards the responses to clients 109.

Server 106 includes non-volatile memory 110, working memory 112, server mass data storage 114, a processing unit 116, and one or more user input/output (I/O) devices 118, all intercommunicating via a server bus 120 (e.g., PCI bus). Non-volatile memory 110 (e.g., read-only memory and/or one or more hard-disk drives) provides storage for data and code which is retained even when server 106 is powered down. Working memory 112 (e.g., random access memory) provides operational memory for server 106, and includes executable code (e.g., an operating system) which is loaded into working memory 112 during start-up. Among other programs, working memory 112 includes server applications 121 and a communication protocol stack 122. Server applications 121 include network software applications (e.g., FTP, HTTP, etc.) which allow server 106 to function as a network server. Communications protocol stack 122 is a standard protocol stack (e.g., TCP/IP) which facilitates communication with other machines over an internetwork. Standard protocol stacks are well known in the art. See, for example, W. Richard Stevens, *TCP/IP Illustrated, Vol. 1* (Addison-Wesley, 1994), which is incorporated herein by reference. Server mass data storage 114 provides data storage (e.g., one or more hard disk drives) for data (e.g., HTML pages, graphics files, etc.), which the server provides to clients 109(1-n) attached to internetwork 102. Processing unit 116 executes the instructions in working memory 112 to cause server 106 to carry out its primary function (e.g., providing data to and receiving data from clients). I/O devices 118 typically include a keyboard, a monitor, and/or such other devices which facilitate user interaction with server 106. Each of the above described components is typically found in a network server such as an internet web server.

Adapter card 108 includes non-volatile memory 123, working memory 124, a processing unit 126, a bus protocol bridge 128, and a network controller 129, all intercommunicating via an adapter bus 130. Non-volatile memory 123 provides storage for data and code (e.g., boot code) which is retained even when adapter 108 is powered down. Processing unit 126 imparts functionality to adapter card 108 by executing the code present in working memory 124. Bus protocol bridge 128 provides an interface between adapter bus 130 and server bus 120, and network controller 129 provides an interface between adapter bus 130 and network media 104.

Working memory 124 provides operational memory for adapter 108, and includes a proxy application 132 and a communication protocol stack 134. Proxy 132 and protocol

stack 134 are loaded from non-volatile memory 123 into working memory 124 at start-up. Optionally, proxy 132 and protocol stack 134 can be loaded from one or more alternative sources, including but not limited to non-volatile memory 110 or server mass data storage 114 of server 106. Proxy 132, when executed by processing unit 126, establishes and manages the 5 above described connections between adapter 108 and server 106 and between adapter 108 and clients 109.

In this particular embodiment of the invention, protocol stacks 122 and 134 are standard (e.g., TCP/IP) protocol stacks. Employing a standard communication protocol stack in adapter 108 facilitates the use of the standard communication software (e.g., protocol stack 10 122) already present in the vast majority of network servers. Those skilled in the art will recognize, however, that this particular element (as well as other described elements, even if not explicitly stated) is not an essential element of the present invention. For example, the present invention may be practiced with custom communication software (e.g., direct 15 communication between server applications 121 and either protocol stack 134 or proxy 132) in both server 106 and adapter 108. Further, in particular embodiments of the invention, this element may be omitted by providing proxy 132 with direct access to the resources (e.g., server mass data storage 114) of server 106.

Adapter card 108 is coupled to server 106 via a bus connection 136 between bus 20 protocol bridge 126 and server bus 120. In this particular embodiment, bus connection 136 is a typical bus expansion slot, for example a PCI slot. Those skilled in the art will recognize, however, that the present invention may be implemented with other types of bus connections, including but not limited to an ISA slot, a USB port, a serial port, or a parallel port. Bus connection 136 facilitates high speed, large packet size, relatively error free (as compared to network connections) communication between proxy 132 and server applications 121, greatly 25 reducing the connection management burden on processing unit 116 of server 106. In summary, proxy 132 (running on processing unit 116) communicates with clients 109 over slow, error prone network connections, and then communicates with server applications 121 on behalf of clients 109 over high speed bus connection 136.

FIG. 2 is a block diagram of working memory 124 showing proxy 132 and protocol 30 stack 134 in greater detail. Those skilled in the art will recognize that while the various software modules of proxy 132 are shown as interconnected functional blocks, the software

modules are actually blocks of executable code stored in working memory 124 that can communicate with one another when executed by processing unit 126 (FIG. 1).

Proxy 132 includes a master process module 202, a plurality of client process modules 204(1-n), a data buffer 206, and an application proxies module 208. Master process module 5 provides overall control and coordination of the various modules of proxy 132. Responsive to a connection request from a client 109 on internetwork 102 (FIG. 1) master process 202 accepts the connection request, initializes a data structure for that client connection in data buffer 206, initiates a new, separate client connection process 204 to handle the connection, and then notifies application proxies 208 that the particular client connection has been 10 established. Each client process 204 handles one such client connection. Application proxies 208 establish and manage bus connections with server 106. Data buffer 206 provides storage for data received from clients 109 and destined for server 106, for data received from server 106 and destined for clients 109, and for connection data relating to established client and/or server connections.

15 Communications protocol stack 134 is a TCP/IP stack including a sockets layer 210, a TCP layer 212, an IP layer 214, and a device layer including a network driver 216 and a server bus driver 218. The functionality of each of the individual layers of protocol stack 134 is well known in the art, and will not, therefore, be discussed in detail herein. Connections between the various modules of proxy 132 and server applications 121 are established 20 through sockets layer 210, TCP layer 212, IP layer 214 and server bus driver 218. Connections between the various modules of proxy 132 are established with clients 109 through sockets layer 210, TCP layer 212, IP layer 214 and network driver 216.

FIG. 3 is a block diagram showing application proxies module 208 to include a plurality of application specific proxies 208(1-f), including a hypertext transfer protocol (HTTP) proxy 208(1), a pass-through proxy 208(2), a security proxy 208(3), and an “other” proxy 208(f). Master process 202 notifies application proxies 208 of an established client connection, by configuring one or more of the application specific proxies 208(1-f) to service the client connection. One means of configuring an application specific proxy (e.g., HTTP proxy 208(1)) is to enter a client process identifier in the processing queue of the application 25 specific proxy. 30

Master process 202 determines which of the application specific proxies to implement for a particular client process from the port number included in the client connection request. It is standard practice to use well known port numbers to identify particular network applications and/or protocols (e.g., file transfer protocol (FTP), HTTP, etc.). For example, 5 port number 80 corresponds to an HTTP connection request. Master process 202 therefore notifies HTTP proxy 208(1) of all client process' 204 initiated in response to a connection request indicating port 80.

HTTP proxy 208(1) monitors each of the client processes of which it is notified. When HTTP proxy 208(1) determines that a complete HTTP request is received and stored in 10 data buffer 206 by a client process (e.g., 204(n)), HTTP proxy 208(1) opens a connection to the server, transmits the request to the server, receives a response from the server, stores the response in data buffer 206 and then closes the server connection. The server response is then transmitted to client 109(n) by the associated client process 204(n).

When master process 202 receives a connection request with a port number that does 15 not correspond to any of the other application specific proxies, master process 202 notifies pass-through proxy 208(2). Pass-through proxy 208(2) simply opens a server connection, transfers the data received from the associated client process 204 from data buffer 206 to server 106, and then closes the server connection.

Master process 202 may notify some application specific proxies of all client 20 connections, regardless of the associated port number. For example, security proxy 208(3) is operative to screen all client connection requests by, for example, terminating any client process initiated in response to a connection request lacking some indicia of authorization, prior to implementing one of the other application specific proxies.

“Other” proxy 208(f) is included in FIG. 3 to show that application proxies 208 can 25 include any currently known or future developed proxy that is desirable for a particular application, including but not limited to, caching HTTP proxy applications, electronic mail applications, and file transfer applications.

FIG. 4 shows an example of client data structures 402(1-n) and proxy data structures 404(1-f), implemented in data buffer 206 to effect the transfer of data through proxy 132. 30 Master process 202 creates and initializes one client data structure 402 for each client process

204, and one proxy data structure 404 for each application specific proxy in application proxies 208.

Each client data structure 402 includes a client socket 406, a server socket 408, a connection state 410, an input queue 412, an output queue 414, and application proxy data 416. For each client connection (e.g., connection (n)), client socket 406(n) and server socket 408(n) each include the IP address and port number of the client 109(n) and server 106, respectively, thus uniquely associating each client data structure 402(n) with a single one of client processes 204(n). Connection state 410(n) indicates the current status (e.g., complete request received, response received, etc.) of the connection (n). Input queue 412(n) is used to store and accumulate data received from client 109(n) by the client process 204(n) associated with the particular data structure 402(n). Output queue 414(n) is used to store data from application proxies 208 which is to be forwarded to client 109(n) by client process 204(n). Application proxy data 416(n) is provided to store any information specific to a particular application proxy (e.g., flags, etc.).

Each proxy data structure(e.g., 404(f)) includes a client queue 418(f), a client ready queue 420(f), and a read pending queue 422(f). Client queue 418(f) includes a client process descriptor (e.g., a pointer to a related client data structure 402) for each client process 204 associated with the particular application proxy (f) to which the proxy data structure 404(f) corresponds. Client ready queue 420(f) includes a client process descriptor for each client data structure 402 that has data in its input queue 412 that is ready to be processed (e.g., transferred to server 106) by the associated application proxy (f). Read pending queue 422(f) includes the client process descriptor for each client process that is awaiting a response from server 106.

Those skilled in the art will understand that the above described client data structure 402 and proxy data structure 404 are exemplary in nature, and that other data structures may be employed with the present invention. The configuration of such alternate data structures will necessarily depend on the function and structure of the particular application specific proxies that are employed.

FIG. 5 is a flowchart summarizing a particular method 500 of managing connections between clients and a server according to the present invention. In a first step 502 proxy 132 establishes a network connection with a client 109, and then in a second step 504 receives a

communication (e.g., an HTTP request) from client 109 via the network connection. Next, in a third step 506, proxy 132 establishes a bus connection with server 106, and then in a fourth step 508 forwards the received client communication to server 106 via the bus connection.

Then, in a fifth step 510, proxy 132 receives a response (e.g., HTML data) to the client  
5 communication from server 106, and in sixth step 512 transmits the response to client 109 via the client network connection. Finally, in a seventh step 514, proxy 132 determines whether there is a signal to terminate (e.g., shut down), and if there is a signal to terminate, then method 500 ends. If there is no signal to terminate in seventh step 514, then method 500 returns to first step 502, to establish a network connection with another client 109.

10 FIG. 6 is a flow chart summarizing one particular method 600 of performing first step 502 (establishing a network connection with a client) of method 500. In a first step, master process 202 connects to internetwork 102. Then in a second step 604 master process 202 listens to the traffic on internetwork 102 to determine if there is a connection request from a client 109. If there is no client connection request, then method 600 ends. If there is a  
15 connection request from a client 109(n), then in a third step 606, master process 202 accepts the connection request from client 109(n), initiates a client process 204(n) to handle the connection, and initializes a client data structure 402(n) in data buffer 206. Next, in a fourth step 608, master process 202 discerns a proxy application identifier (e.g., a port number) from the client connection request and notifies one or more of application proxies 208(1-f),  
20 depending on the value of the identifier, by writing the client process descriptor (e.g., pointer to client data structure 402(n)) into the client queues 418 of the respective proxy data structures 404. Finally, in a fifth step 610, master process 202 determines whether the maximum allowed number of client connections are open. If the maximum number of client connections are open, then method 600 ends. If the maximum number of client connections  
25 are not open, then method 600 returns to second step 604, to listen for another connection request.

FIG. 7 is a flow chart summarizing a method 700 of performing second step 504 (receiving a communication from a client 109) of method 500. In a first step 702, master process 202 determines whether there are any client processes 204 to be processed to receive data. If master process 202 has already processed all of client processes 204(1-n), then method 700 ends. If not, then in a second step 704, master process 202 calls the first client

process 204(1). Then, in a third step 706, client process 204(1) checks its client connection (e.g., the TCP buffer) to determine whether there is any data coming in from client 109(1). If there is no incoming data for the first client process 204(1), then method 700 returns to first step 702 to process any remaining client processes 204(2-n). If, in third step 706, client  
5 process 204(1) determines that there is incoming data from client 109(1), then in a fourth step 708, client process 204(1) checks client data structure 402(1) to determine if input queue 412(1) is available to receive client data. If input queue 412(1) is not available, then method 700 returns to first step 702 to process any remaining client processes 204(2-n). If in fourth step 708, client process 204(1) determines that input queue 412(1) is available to receive data,  
10 then in a fifth step 710, client process 204(1) transfers the incoming client data into input queue 412(1). Then, in a sixth step 712, client process 204(1) determines whether the data accumulated in input queue constitutes a complete request (i.e., data ready to be transferred to server 106, for example a complete HTTP request). If the data does not constitute a complete request, then method 700 returns to first step 702 to process any remaining client processes  
15 204(2-n). If, however, client process 204(1) determines in sixth step 712 that the data in input queue 412(1) constitutes a complete request, then, in a seventh step 714, client process notifies proxy applications 208 that there is a complete request by, for example, setting connection state 410(1) to so indicate. Then, method 700 returns to first step 702 to determine whether there are any more client processes 204(2-n) to process.

20 FIG. 8 is a flow chart summarizing a method 800 of performing third step 506 (establishing a bus connection with server 106) of method 500. In a first step 802, a first one of application proxies 208(1) retrieves the first client descriptor from client queue 418(1) of its proxy data structure 404(1). Then in a second step 804, proxy 208(1) checks the connection state 412 of the client data structure 402 identified by the first client descriptor to determine if the first client has a complete request in its input queue 412. If connection state 412 indicates a complete request, then in a third step 806, proxy 208(1) adds the client descriptor to its client ready queue 420(1). Next, in a fourth step 808, proxy 208(1) determines whether the maximum number of connections to server 106 are open. If the maximum number of server connections are already open, then method 800 ends. If the  
25 maximum number of server connections are not already opened, then in a fifth step 810, proxy 208(1) opens a bus connection with server 106 and writes connection information in  
30

server socket 408 of the associated client data structure 402. Next, in a sixth step 812, proxy 208(1) determines whether it has checked the last client descriptor in its client queue 418(1). If the last descriptor has been checked, then method 800 ends. Otherwise, method 800 returns to first step 802 to retrieve the next descriptor in client queue 418(1). If, in second 5 step 804, proxy 208(1) determines that a complete client request has not been received, then method 800 proceeds directly to sixth step 812. Once all of the descriptors in client queue 418(1) of proxy data structure 404(1) have been processed, method 800, or a similar method, is repeated for each of the other application proxies 208(2-f).

FIG. 9 is a flow chart summarizing a method 900 of performing fourth step 508

- 10 (forwarding a client communication to server 106) of method 500. In a first step 902, proxy 208(1) retrieves the first client descriptor from the client ready queue 420(1) of its proxy data structure 404(1). Then, in a second step 904, proxy 208(1) checks the server socket 408 of the first client's data structure 402 to determine whether a server connection is open. If a server connection is open, then in a third step 906, proxy 208(1) transfers the client data (e.g., 15 HTTP request) from the client input queue 412 to server 106 over the open server connection. Next, in a fourth step 908, proxy 208(1) moves the client descriptor from the client ready queue 420(1) to the read pending queue 422(1). Then in a fifth step 910, proxy 208(1) determines whether the last client in the client ready queue 420(1) has been checked. If not, then method 900 returns to first step 902 to check the next client in client ready queue 420(1). 20 If the last client has been checked, then method 900 ends. If, in second step 904, proxy 208(1) determines that there is no server connection open for a particular client, then method 900 proceeds directly to fifth step 910 to determine whether the last client in the client ready queue 420(1) has been checked. Once all of the descriptors in client ready queue 420(1) of proxy data structure 404(1) have been processed, method 900, or a similar method, is 25 repeated for each of the other application proxies 208(2-f).

FIG. 10 is a flow chart summarizing a method 1000 of performing fifth step 510 (receive a response from server 106) of method 500. In a first step 1002, proxy 208(1) determines whether read pending queue 422(1) is empty, and if so method 1000 ends. If read pending queue 422(1) is not empty, then in a second step 1004 proxy 208(1) retrieves the first 30 client descriptor from the read pending queue. Next, in a third step 1006, proxy 208(1) checks the open server connection identified in server socket 408 of the client data structure

402 identified by the first client descriptor, to determine whether there is any incoming server data (i.e., a response to the client request) on that connection. If there is no incoming server data on that connection, then method 1000 returns to second step 1004 to check the next client in the read pending queue. If there is incoming server data, then in a fourth step 1008, 5 proxy 208(1) checks to determine whether output queue 414 of the client data structure 402 identified by the client descriptor is available. If output queue 414 is not available, then method 1000 returns to second step 1004 to check the next client descriptor in read pending queue 422(1). If output queue 414 is available, then, in a fifth step 1010, proxy 208(1) moves the incoming server data into the output queue 414 of client data structure 402. Next, in a 10 sixth step 1012, proxy 208(1) determines whether the server data includes an “end of file” indicator. If not, then in a seventh step 1014, proxy 208(1) checks to determine whether the last client descriptor in read pending queue 422(1) has been processed. If so, method 1000 ends. If not, method 1000 returns to step 1004 to read the next client descriptor from read pending queue 422(1).

15       If, in sixth step 1012, proxy 208(1) determines that the server data includes an end-of-file indicator, then method 1000 proceeds to an eighth step 1016, wherein proxy 208(1) removes the client descriptor from the read pending queue, and then in a ninth step 1018 closes the server connection. After ninth step 1018, method 1000 returns to seventh step 1014. Once all of the descriptors in read pending queue 422(1) of proxy data structure 404(1) 20 have been processed, method 1000, or a similar method, is repeated for each of the other application proxies 208(2-f).

FIG. 11 is a flow chart summarizing a method 1100 of performing sixth step 512 (transmitting data to clients 109) of method 500. In a first step 1102, master process 202 calls first client process 204(1). Then, in a second step 1104, first client process 204(1) determines 25 whether there is any data in output queue 414(1) of client data structure 402(1). If there is no data in output queue 414(1), then method 1100 returns to first step 1102 where master process 202 calls the next of the remaining client processes 204(2-n). If, however, in second step 1104, client process 204(1) determines that there is data in output queue 414(1), then in a third step 1106, client process 204(1) determines whether the client network connection is 30 ready to receive data. If the client network connection is not ready, then method 1100 returns to first step 1102. If the client network connection is ready, then in a fourth step 1108, client

process 204(1) moves at least a portion of the data in output queue 414(1) to the client connection (e.g., the TCP output buffer). Next, in a fifth step 1110, master process 202 determines whether the last client process has been called. If so, method 1100 ends. If not, method 1100 returns to first step 1102 to call the next of the remaining client processes 203(2-n).

FIG. 12 is a block diagram showing an alternate client data structure that can be implemented in data buffer 206 to significantly increase the number of simultaneous client connections that can be maintained by proxy 132. In particular, the number of simultaneous client connections that can be maintained depends on the memory capacity of data buffer 206 available to buffer the data flowing between clients 109 and server 106. Some applications (e.g., “buddy list” applications) open client connections with a server, but rarely exchange data with the server. Other applications (e.g., web browsers) open a client connection with a server and request data (e.g., a web page) from a server, but do not close the connection after the data is received. At any given time, therefore, a large number of client connections may be open, and tying up buffer space, but not actively exchanging data with the server.

Dissociating the allocation of data buffer space from the client data structure substantially increases the number of simultaneous client connections that can be open, because valuable memory space is not wasted by such idle client connections.

Dissociated client data structures 1202(1-n) are similar to client data structures 402(1-n), except that input queues 412(1-n) and output queues 414(1-n) are replaced with input buffer identifiers 1204(1-n) and output buffer identifiers 1206(1-n), respectively. Input buffer identifiers 1204 and output buffer identifiers 1206 store pointers that link the client data structures 1202 to buffers that are allocated to the particular client process 204 associated with the client data structure 1202. When master process module 202 forks client data structures 1202(1-n), default values (e.g., all bits set to “1”) are loaded into input buffer identifier 1204 and output buffer identifier 1206 to indicate that no data buffers have been allocated. Storage for the data that passes between clients 109 and server 106 is provided by a buffer pool 1208 that includes a buffer status information section 1210 and a plurality of general-purpose buffers 1212(0-Z). Buffer status section 1210 stores information relating to the allocation and use of buffers 1212(0-Z), as will be described in greater detail below. Each of buffers 1212(0-Z) is a general-purpose buffer of predetermined size (e.g., 2KB). The

terminal storage locations 1214(0-Z) of each of buffers 1212(0-Z) are reserved for pointers that facilitate the linking of two or more of buffers 1212(0-Z) to create a buffer of greater capacity.

FIG. 13 is a block diagram showing buffer status information section 1210 in greater detail. Buffer status section 1210 includes a plurality of registers 1302(0-Z), each associated with a corresponding one of buffers 1212(0-Z). Each of registers 1302(0-Z) includes a start address storage location 1304(0-Z), a length of valid data storage location 1306(0-Z), and a status flag 1308(0-Z), respectively. Start address storage location 1304 stores a pointer to the beginning of the buffer 1212 associated with the particular register 1302. Length of valid data storage location 1306 stores two values, one that indicates how much data has been written into associated buffer 1212 (data written value), and one that indicates how much data has been transferred out of buffer 1212 (data read value). Status flag 1308 indicates whether or not associated buffer 1212 has been allocated to a client data structure 1202 or is available for allocation.

FIG. 14 is a flow chart summarizing an alternate method 1400 of performing second step 504 (receive request via client connection) of method 500, that utilizes client data structures 1202(1-n) and buffer pool 1208. In a first step 1402, master process 202 determines whether there are any client processes 204 to be processed to receive data. If master process 202 has already processed all of client processes 204(1-n), then method 1400 ends. If not, then in a second step 1404, master process 202 calls the first client process 204(1). Then, in a third step 1406, client process 204(1) checks its client connection (e.g., the TCP buffer) to determine whether there is a data request coming in from client 109(1). If there is no incoming data for the first client process 204(1), then method 1400 returns to first step 1402 to process any remaining client processes 204(2-n). If, in third step 1406, client process 204(1) determines that there is an incoming data request from client 109(1), then in a fourth step 1408, client process 204(1) checks input buffer identifier 1204(1) to determine whether one of buffers 1212(0-z) has been allocated as an input buffer to client process 204(1). If there is an input buffer allocated, then in a fifth step 1410, client process 204(1) moves the client data into the allocated input buffer. Then, in a sixth step 1412, client process 204(1) determines whether the data accumulated in the allocated input buffer constitutes a complete request (i.e., data ready to be transferred to server 106, for example a

complete HTTP request). If the data does not constitute a complete request, then method 1400 returns to first step 1402 to process any remaining client processes 204(2-n). If, however, client process 204(1) determines in sixth step 1412 that the data in the input buffer constitutes a complete request, then, in a seventh step 1414, client process 204(1) notifies proxy applications 208 that there is a complete request by, for example, setting connection state 1210(1) to so indicate. Then, method 1400 returns to first step 1402 to determine whether there are any more client processes 204(2-n) to process.

If in fourth step 1408, client process 204 determines that no input buffer has been allocated to client process 204(1), then in an eighth step 1416, client process 204(1) examines status flags 1308 of buffer status information 1210 to determine whether one of buffers 1212(0-Z) is available. If one of buffers 1212(0-Z) is available, then in a ninth step 1418, client process 204(1) allocates the available buffer 1212 to itself, by writing a pointer to the available buffer's status register 1302 into input buffer identifier 1204(1) of client data structure 1202(1), and by setting status flag 1308 in the buffer's register 1302 to indicate that the buffer has been allocated. Next, method 1400 proceeds to fifth step 1410, where client process 204(1) moves the data request into the allocated input buffer 1212.

Method 1400 is similar to method 700, except that none of the buffer resources are allocated as input buffers until it is determined that a client communication includes at least a portion of a data request. A large number of idle client connections can therefore be established and maintained without using any of the valuable buffer space. The buffer space is thus available to support the management of a substantially increased number of active (exchanging data) connections by proxy 132.

As used herein, the term "data request" is intended to mean any data sent by a client 109 and bound for server 106. The term data request should be interpreted broadly to include requests from a client 109 for data from server 106 (e.g., file download request), as well as submissions of data from a client 109 to server 106 (e.g., file upload). The term "data request" does not, however, include connection management data (e.g., connection requests). Thus, according to method 1400, client processes 204(1-n) only allocate buffers 1212(0-Z) when they are needed, as opposed to every time a client connection is opened.

FIG. 15 is a flow chart summarizing an alternate method 1500 of performing fourth step 508 (forward client request to server via bus connection) of method 500. In a first step

1502, proxy 208(1) retrieves the first client descriptor from the client ready queue 420(1) of its proxy data structure 404(1). Then, in a second step 1504, proxy 208(1) checks the server socket 1209 of the first client's data structure 1202 to determine whether a server connection is open. If a server connection is open, then in a third step 1506, proxy 208(1) checks the  
5 output buffer identifier 1206 of the first client's data structure 1202 to determine whether one of buffers 1212(0-Z) has been allocated to the first client's data structure as an output buffer. If proxy 208(1) determines that an output buffer has been allocated, then in a fourth step 1508 proxy 208(1) transfers the client data request (e.g., HTTP request) from the input buffer allocated to the first client to server 106 over the open server connection. Next, in a fifth step  
10 1510, proxy 208(1) moves the client descriptor from the client ready queue 420(1) to the read pending queue 422(1). Then in a sixth step 1512, proxy 208(1) determines whether the last client in the client ready queue 420(1) has been checked. If not, then method 1500 returns to first step 1502 to check the next client in client ready queue 420(1). If the last client has been checked, then method 1500 ends. If, in second step 1504, proxy 208(1) determines that there  
15 is no server connection open for a particular client, then method 1500 proceeds directly to sixth step 1510 to determine whether the last client in the client ready queue 420(1) has been checked. Once all of the descriptors in client ready queue 420(1) of proxy data structure 404(1) have been processed, method 1500, or a similar method, is repeated for each of the other application proxies 208(2-f).

20 If in third step 1506, proxy 208(1) determines that no output buffer has been allocated to the first client's data structure 1202, then in a seventh step 1514, proxy 208(1) searches the buffer status information section 1210 of buffer pool 1208 to determine if any of buffers 1212(0-Z) are available for allocation. If one of buffers 1212(0-Z) are available for allocation, then in an eighth step 1516 proxy 208(1) allocates the available one of buffers  
25 1212(0-Z) to the first client's data structure 1202 by writing a pointer to the available buffer's status register 1302 into input buffer identifier 1204 of the first client's data structure 1202, and by setting the status flag 1308 in the buffer's status register 1302 to indicate that the buffer has been allocated. Then, method 1500 proceeds to fourth step 1508. If none of buffers 1212(0-Z) are available, then method 1500 returns to first step 1502 to process the  
30 next client in client ready queue 420(1).

Method 1500 is similar to method 900, except that buffers 1212(0-Z) are not allocated as output buffers (for storing server responses) until after a server connection is open. The particular triggering event for allocating an output buffer is not, however, considered to be an essential element of the present invention. For example, output buffers can be allocated sooner (e.g., when proxy 208 determines that a complete client request has been received in second step 804 of method 800 (FIG. 8)), or they can be allocated later (e.g., after proxy 208 determines that there is server data available in third step 1006 of method 1000 (FIG. 10)). As long as the buffers are allocated at a time when there is an increased probability that they will be used shortly after allocation, buffer efficiency is increased.

10 FIG. 16 is a flow chart summarizing a method 1600 for writing data to an allocated input or output buffer. Method 1600 will be described with reference to writing data to an allocated input buffer, but is equally well suited to writing server data to an output buffer. In a first step 1602, a client process (e.g., client process 204(1)) uses input buffer identifier 1204(1) to retrieve the buffer status information (the start address 1304 and the length of valid data 1306) for the allocated buffer 1212. Then, in a second step 1604, client process 204(1) transfers a first block of the available client data into the allocated buffer 1212. Client process 204(1) calculates the storage address for the block of data by adding the length of valid data 1306 (data written value) to the start address 1302 of the buffer. Then, in a third step 1606, client process 204(1) updates the buffer status information by incrementing the length of valid data 1306 (data written value) by the size of the data block written to the allocated buffer 1212. Next, in a fourth step 1608, client process 204(1) determines whether the transferred block of data included an end-of-data indicator, and if so then method 1600 ends.

25 If, in fourth step 1608, client process 204(1) determines that the transferred data block did not include an end-of-file indicator, then in a fifth step 1610 client process 204(1) determines whether the allocated buffer is full by comparing the updated length of valid data 1306 to the known size of buffer 1212. If the data buffer 1212 is not full, then method 1600 returns to second step 1604 to transfer the next block of data. If, however, the allocated data buffer 1212 is full, then in a sixth step 1612 client process 204(1) searches buffer status 30 information 1210 to determine whether any of buffers 1212(0-Z) are available. If there are no unallocated buffers in buffer pool 1208, then method 1600 ends, but if client process 204(1)

finds an unallocated buffer in buffer pool 1208, then in an eighth step 1616 client process 204(1) links the unallocated buffer to the previously allocated input buffer. Client process 204(1) links the buffers by storing a pointer to the status register 1302 of the second allocated buffer in the last position 1214 of the first buffer, updating the length of valid data 1306, and 5 setting the status flag 1308 in the status register 1302 of the new buffer to indicate that the new buffer is allocated. After linking the input buffers in eighth step 1616, method 1600 returns to second step 1604 to transfer the next block of data. Method 1600 continues until either all of the available data is written to the input buffer, or until all of buffers 1212(0-Z) are allocated.

10 FIG. 17 is a flow chart summarizing a method 1700 of transferring data out of an allocated input or output buffer. Method 1700 is described with reference to transferring server data out of an allocated output buffer, but is equally well suited for transferring client data out of an input buffer. In a first step 1702, a client process 204 uses output buffer identifier 1206(1) to retrieve the buffer status information for the allocated output buffer 1212. Then, in a second step 1704, client process 204 transfers a first block of the stored 15 server data out of the allocated output buffer 1212 to the client 109. In a third step 1706, client process 204 updates the buffer status information 1210 by incrementing the data read value of valid data 1306 by the size of the data block transferred. Next, in a fourth step 1708 client process 204 determines whether the output buffer is empty, by comparing the data 20 written and data read values of the length of valid data 1306. If the data read value is equal to the data written value or the size of buffer 1212, then buffer 1212 is empty. If buffer 1212 is empty, then in a fifth step 1710 client process 204 determines whether the empty buffer is linked to any additional buffers 1212. If the data written value of the length of valid data 1306 is smaller than the known size of buffers 1212, then there are no linked buffers. If there 25 are no linked buffers, then in a sixth step 1712, client process 204 releases buffer 1212, by changing its status flag 1308 to indicate that it is unallocated and resetting the length of valid data values 1306, and then method 1700 ends.

If in fifth step 1710 client process 204 determines that the empty buffer is linked to an additional buffer, then in a seventh step 1714 client process 204 unlinks the empty buffer 30 from any buffers still containing data as follows. First, client process 204 copies the pointer from the terminal location 1214 of the empty buffer into the output buffer identifier 1206 of

the client data structure 1202. The copied pointer points to the status register 1302 of the next linked buffer, instead of the status register of the emptied buffer. Next, client process 204 copies the length of valid data values 1306 from the status register 1302 of the empty buffer into the status register 1302 of the next linked buffer, and then decrements both the data  
5 written and the data read values by an amount equal to the capacity of one buffer 1212. After unlinking the buffers, client process 204 releases the empty buffer, in an eighth step 1716, as described above with respect to sixth step 1712. After the empty buffer is released in eighth step 1716, method 1700 returns to second step 1704 to begin transferring data out of the next linked buffer.

10 The description of particular embodiments of the present invention is now complete. Many of the described features may be substituted, altered or omitted without departing from the scope of the invention. For example, the operative components of adapter 108 (e.g., processing unit 126 and proxy 132) can be incorporated directly into a server instead of being provided in a removable adapter card. Further, alternate data structures may be substituted  
15 for the exemplary data structures provided. Additionally, the particular orders of methods and routines disclosed herein are not considered to be essential elements of the present invention. As yet another example, master process 202 can be configured to open a predetermined number of persistent bus connections with server 106 at start-up, and manage the use of those connections by application proxies 208(1-f), thus eliminating the need for  
20 server 106 to repetitively open and close the bus connections. These and other deviations from the particular embodiments shown will be apparent to those skilled in the art, particularly in view of the foregoing disclosure.

We claim:

1. A method for managing connections between at least one client and a server, said method comprising:

- 5        establishing a network connection with one of said clients via a network;  
receiving a communication from said client via said network connection;  
determining whether said communication includes a data request;  
allocating an input buffer to said client connection if said communication includes a  
data request;
- 10      storing said data request in said allocated buffer space;  
establishing a bus connection with said server via an internal bus of said server; and  
forwarding said data request to said server via said bus connection.

2. A method according to Claim 1, wherein said step of establishing a bus connection  
15 with said server includes allocating an output buffer to store said server's response to said  
data request when received.

3. A method according to Claim 2, wherein said step of forwarding said data request  
to said server includes releasing said input buffer.

20      4. A method according to Claim 3, further comprising:  
receiving a response to said data request from said server; and  
storing said data request in said output buffer.

25      5. A method according to Claim 4, wherein said step of storing said response in said  
output buffer comprises:  
determining whether said output buffer is full; and  
if said output buffer is full, allocating additional buffer space to said output buffer.

6. A method according to Claim 5, wherein said step of allocating additional buffer space to said output buffer includes linking said output buffer to at least one other buffer from a pool of available buffers.

5        7. A method according to Claim 6, further comprising:  
forwarding said server response to said client via said network connection; and  
releasing each said linked buffer as it is emptied.

10      8. A method according to Claim 5, wherein said step of allocating additional buffer space to said output buffer is repeated until the entire server response is received and stored.

9. A method according to Claim 4, further comprising:  
forwarding said server response to said client via said network connection; and  
releasing said output buffer.

15      10. A method according to Claim 1, wherein:  
said step of allocating an input buffer includes determining whether an input buffer has been previously allocated to the client connection; and  
if an input buffer has been previously allocated, not allocating another input buffer.

20      11. A method according to Claim 10, wherein step of storing said data request in said allocated input buffer includes:  
determining whether said input buffer is full; and  
if said input buffer is full, allocating additional buffer space to said input buffer.

25      12. A method according to Claim 11, wherein said step of allocating additional buffer space to said input buffer is repeated until the entire data request is received and stored.

13. A method according to Claim 11, wherein said step of allocating additional buffer space to said input buffer includes linking said input buffer to at least one other buffer from a pool of available buffers.

5        14. A method according to Claim 13, wherein said step of forwarding said data request to said server includes releasing each of said linked buffers as they are emptied.

15. A method for managing connections between at least one client and a server, said method comprising:

10        establishing a network connection with one of said clients via a network;  
receiving a communication from said client via said network connection;  
determining whether a complete data request has been received from said client;  
allocating an output buffer to store a response from said server, only after a complete  
data request has been received;  
15        establishing a bus connection with said server via an internal bus of said server; and  
forwarding said data request to said server via said bus connection.

16. A method according to Claim 15, further comprising:  
receiving a response to said data request from said server; and  
20        storing said data request in said output buffer.

17. A method according to Claim 16, wherein said step of storing said response in said output buffer comprises:

25        determining whether said output buffer is full; and  
if said output buffer is full, allocating additional buffer space to said output buffer.

18. A method according to Claim 17, wherein said step of allocating additional buffer space to said output buffer includes linking said output buffer to at least one other buffer from  
30        a pool of available buffers.

19. A method according to Claim 18, further comprising:  
forwarding said server response to said client via said network connection; and  
releasing each said linked buffer as it is emptied.
- 5 20. A method according to Claim 17, wherein said step of allocating additional buffer space to said output buffer is repeated until the entire server response is received and stored.
- 10 21. A method according to Claim 16, further comprising:  
forwarding said server response to said client via said network connection; and  
releasing said output buffer.
22. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 1.
- 15 23. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 2.
24. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 3.
- 20 25. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 4.
26. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 5.
27. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 6.
- 30 28. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 7.

29. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 8.

5       30. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 9.

31. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 10.

10      32. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 11.

15      33. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 12.

34. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 13.

20      35. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 14.

36. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 15.

25      37. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 16.

30      38. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 17.

39. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 18.

40. A computer readable medium having code embodied therein for causing an 5 electronic device to perform the steps of Claim 19.

41. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 20.

10 42. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 21.

43. An adapter card for coupling a server with an internal bus to a network, said adapter card comprising:

15 a network controller for communicating with clients on said network;  
a memory device for storing data and code, said memory device including a plurality  
of buffers, and said code including a proxy application responsive to  
communications from said clients and operative to allocate an input buffer to a  
client connection only if a communication received via said client connection  
20 includes a data request;  
a processing unit coupled to said memory device for executing said code; and  
a protocol adapter coupled to said processing unit, and adapted to couple to said  
internal bus, for communicating with said server.

25 44. An adapter card according to Claim 43, wherein said proxy application is further  
operative to allocate an output buffer to said client connection only after receiving a complete  
data request via said client connection, said output buffer for storing a response to said data  
request from said server.

45. An adapter card according to Claim 44, wherein said proxy application is further operative to forward said complete data request to said server, and to release said input buffer thereafter.

5       46. An adapter card according to Claim 45, wherein said proxy application, responsive to receipt of a response to said data request from said server, is further operative to:

store said response in said output buffer;

forward said response to a client via said client connection; and

10      release said output buffer after said response has been forwarded to said client.

47. An adapter card according to Claim 44, wherein each of said plurality of buffers is available to said proxy application for use as said input or said output buffer.

15      48. An adapter card according to Claim 43, wherein said proxy application is further operative to forward said data request to said server, and to release said input buffer thereafter.

20      49. An adapter card for coupling a server with an internal bus to a network, said adapter card comprising:  
          a network controller for communicating with clients on said network;  
          a memory device for storing data and code, said memory device including a plurality of buffers, and said code including a proxy application responsive to communications from said clients and operative to allocate an output buffer to a client connection only after a complete data request is received via said client connection;  
          a processing unit coupled to said memory device for executing said code; and  
          a protocol adapter coupled to said processing unit, and adapted to couple to said internal bus, for communicating with said server.

25  
30

50. An adapter card according to Claim 49, wherein said proxy application, responsive to receipt of a response to said data request from said server, is further operative to:

- store said response in said output buffer;
- forward said response to a client via said client connection; and
- release said output buffer after said response has been forwarded to said client.







FIG. 1

**FIG. 2****FIG. 3**






FIG. 5








FIG. 6

FIG. 7

**FIG. 8****FIG. 9**

**FIG. 10**




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17